Citation histories of papers: sometimes the rich get richer, sometimes they don't

Michael J. Hazoglou,[1] Vivek Kulkarni,[2] Steven S. Skiena,[2] and Ken A. Dill[1,3,4]
[1)] *Department of Physics & Astronomy, Stony Brook University*
[2)] *Department of Computer Science, Stony Brook University*
[3)] *Laufer Center for Physical and Quantitative Biology*
[4)] *Department of Chemistry, Stony Brook University*

We describe a simple model of how a publication's citations change over time, based on pure-birth stochastic processes with a linear cumulative advantage effect. The model is applied to citation data from the Physical Review corpus provided by APS. Our model reveals that papers fall into three different clusters: papers that have rapid initial citations and ultimately high impact (*fast-hi*), fast to rise but quick to plateau (*fast-flat*), or late bloomers (*slow-late*), which may either never achieve many citations, or do so many years after publication. In *fast-hi* and *slow-late*, there is a rich-get-richer effect: papers that have many citations accumulate additional citations more rapidly while the *fast-flat* papers do not display this effect. We conclude by showing that only a few years of post-publication statistics are needed to identify high impact (*fast-hi*) papers.

arXiv:1703.04746v1 [cs.DL] 14 Mar 2017

I. INTRODUCTION

Every research publication has a citation history. A paper's citation history contains more information about the paper than any single-number measure alone does, such as its citation count. Here, we develop a generative model to describe the citation histories of publications, and we ask how predictable a paper's future citations are from its initial citations (citation history).

We begin with a few general observations about citations. Like a person, a publication goes through birth-and-death stages. Soon after it is 'born', the paper receives some initial number of citations per year (its citation rate). Later, when interest in the paper has waned, the citation rate diminishes toward zero. Certain papers have high citation rates; others have low rates. Some papers continue receiving citations over a long lifetime, while others die out quickly. Interestingly, one of the earliest papers in the field of bibliometrics, by Lotka in 1926, leading to what is now called 'Lotka's Law[1,2], is itself still quite highly cited.

When looking at multiple papers collected together as a set – for example, the collected papers of an individual or over a field of research – it is found that a plot of the numbers of papers as a function of the citation counts of those papers shows distributions shaped like a power-law[3] or like the exponential of a digamma function[4]. The non-exponential tails of these distributions indicate a 'rich-get-richer' feature: the more citations a paper already has, the higher the rate at which it receives new citations.

These features are captured in a recent model (WSB) of the time-dependent histories of citations over time[5–7]. Here, we go beyond that model in two ways. First, we provide a more microscopic, or generative, model for the underpinnings of citations. Second, we introduce a different but related mechanistic model that we call the *Direct-Indirect* (DI) model[4], which has a simple mechanistic interpretation and captures observed data with fewer unexplained parameters. We use our models to analyze a set of 151,082 papers from the Physical Review Corpus of the American Physical Society across all the disciplines represented by papers published by APS journals. The model is then used to classify papers into particular categories of performance based on the similarity of the papers.

II. THE PURE BIRTH PROCESS AND THE MATHEMATICS OF A PROCESS WITH CUMULATIVE ADVANTAGE

In this section we review a special case of birth-death processes, the pure birth process. Birth-death processes involve an increase in the state variable by one, called birth and a decrease in the state variable by one called death. Here the state variable will be the total number of citations a paper has accumulated. The pure birth process is useful in bibliometrics citations because papers generally don't disappear (retractions are negligibly rare events). A pure birth process resembles a process of compounded interest; the formalism below follows from that. This formalism allows the user a lot of flexibility in describing bibliometric data: as long as the four defining properties of the birth process are true the rate can be refined with different functional forms to model the data. The birth process is described in terms of a conditional probability distribution $p(N(t+\tau) - N(t) = n|N(t) = m) \equiv p_{m,n}(t,\tau)$ for the number of events $N(t)$ at the time t. The initial condition is that $N(0) = 0$. We assume the following properties[8]:

- The process is Markovian depending only on the current state $N(t), t \geq 0$
- $p_{m,1}(t,\delta t) = \lambda_m(t)\delta t + o(\delta t)$ for $m = 0,1,2,3,\ldots$
- $p_{m,n}(t,\delta t) = o(\delta t)$ for $n > 1$, $m = 0,1,2,3,\ldots$

here $\lambda_m(t)$ is the rate which can depend on the total number of events and the time. These conditions can

be used to define a set of differential equations to solve for $p_{m,n}(t,\tau)$[8], but a general solution can be obtained in terms of a recursive integral equation, which is simpler and straight forward to understand. Since cumulative advantage is important in citation dynamics we will focus on processes with rates $\lambda_n(t) = (an+b)\lambda(t)$ with the probability distribution

$$p_{m,n}(t,\tau) = \frac{\Gamma(b/a+m+n)}{n!\Gamma(b/a+m)}[e^{-a\gamma(t,\tau)}]^{\frac{b}{a}+m}[1-e^{-a\gamma(t,\tau)}]^n \tag{1}$$

where $\gamma(t,\tau) = \int_t^{\tau+t}\lambda(u)du$, and $\Gamma(x+1) = x\Gamma(x)$ is the standard gamma function.

The mean number of events in the time interval τ given m events in the time t is given by

$$\langle n(\tau)|m,t\rangle = \left(\frac{b}{a}+m\right)\left[e^{a\gamma(t,\tau)}-1\right]. \tag{2}$$

Eqn. (1) is negative binomial distribution, which is characteristic of a process with cumulative advantage where the frequency of events increases with the number of events. In the context of citations, this would model the discovery of a paper through the references of another more recent paper. The limit as $a \to 0^+$ describes the situation in which there is no cumulative advantage, and it reduces to a Poisson distribution.

It is worth noting that the main results of Shen, Wang, Song, and Barabási[5,6] are derived from substituting into Eqn. (2) with the choice of $\lambda(t)$ being the same functional form as a log-normal distribution. What was not obtained in those papers was a probability distribution for the process. As a result the process has been mislabeled as a "reinforced Poisson process" when in fact Eqn. (1) is a negative binomial distribution[6].

III. DIRECT AND INDIRECT MECHANISM OF CITATION

Here, we consider the Direct-Indirect (DI) mechanism of citations[4]. In the DI mechanism, a paper B receives a *direct citation* from a paper A when the author of A finds B through some direct search over a database of papers. In contrast, a paper B receives an *indirect citation* when the author of A finds paper B in the references of another paper. The difference in these mechanisms is that of being searched and found. A direct search finds the single paper B. An indirect search finds any one of many different papers that cite B. We model the rate $\lambda_n(t)$ at which a paper having n citations is cited at time t as:

$$\lambda_n(t) = \lambda_{\text{indirect}}(t) + \lambda_{\text{direct}}(t) \tag{3}$$

If we say a paper is found selected almost randomly from a body of relevant papers, then

$$\lambda_{\text{direct}}(t) \propto \frac{1}{N(t)} \tag{4}$$

$N(t)$ is the number of relevant papers at time t, and the indirect rate is proportional to the number of citations a paper has received at a given time $n(t)$, so

$$\lambda_{\text{indirect}}(t) \propto \frac{n(t)}{N(t)} \tag{5}$$

since the number of papers increases exponentially with time $N(t) \propto \exp(rt)$[5,9]. Therefore, we obtain:

$$\lambda_n(t) = (an+b)r\exp(-rt). \tag{6}$$

This choice of rate gives $\gamma(t,\tau) = \exp(-rt) - \exp(-r(t+\tau))$, the mean number of citations from Eq. (2), with $t=0$ and $m=0$ as

$$\langle n(\tau)|0,0\rangle = \frac{b}{a}\left[\exp\left(a(1-e^{-r\tau})\right)-1\right] \tag{7}$$

br is easily interpreted as the initial citation and ar is how much the rate increases with each citation the paper receives within a short time ($\tau \ll 1/r$).

Over long time ranges, we can use this to determine the expected number of new citations a paper will receive given some m and t:

$$\lim_{\tau\to\infty}\langle n(\tau)|m,t\rangle = \left(\frac{b}{a}+m\right)\left[\exp\left(ae^{-rt}\right)-1\right] \tag{8}$$

Setting $m=0$ and $t=0$ gives the expected number of total citations

$$\lim_{\tau\to\infty}\langle n(\tau)|0,0\rangle = \frac{b}{a}\left[\exp\left(a\right)-1\right] \tag{9}$$

which for the limiting case of $a=0$ reduces to b. This expression tells us according to this model that the expected total citations grows exponentially with the cumulative advantage parameter a. Even moderately large values of a would result in unreasonably large values unless the parameter b exponentially decreases with a.

The relevant time scale for this process is $1/r$. With our original motivation, this will give us the doubling time from

$$t_{2\times} = \frac{\ln(2)}{r} \tag{10}$$

where a typical doubling time of 13 years would result in a $r \approx 0.05\,\text{year}^{-1}$.

A Taylor series expansion around $\tau = 0$ of $\langle n(\tau)|0,0\rangle$ gives the expected expected initial rate of citation, which will be called the *velocity* $v(0)$.

$$v(0) \equiv \left.\frac{d}{d\tau}\langle n(\tau)|0,0\rangle\right|_{\tau=0} = br \tag{11}$$

IV. CLUSTERING THE MODEL PARAMETERS SHOWS THAT MOST PAPERS FALL INTO ONE OF THREE CLASSES

We apply this model to a database of 151,082 papers from the Physical Review journals. For each paper, we extract the three parameters of the model: b (proportional to the initial citation velocity br), a (a 'rich-get-richer' parameter), and r, a characteristic inverse time scale. We then cluster papers together in terms of those three parameters using the clustering method called DBSCAN.

Figure 1 shows two plots of different parameter pairs, showing that the model clusters papers into three main categories: the *fast and high-impact fast-hi* (shown in red) has a good initial velocity, benefits from a favorable rich-get-richer contribution, and has a long decay time, the *quick to plateau fast-flat* papers (blue) do not benefit from a rich-get-richer component. Finally, *slow-late* papers (green) tend to be long-lived and to benefit from acceleration after a long delay; they are sleepers. Examples of the three corresponding time trajectories are shown in Figure 2.

Figure 3 makes two points. First, it shows that the initial citation rate $v(0)$ and the decay rate r of a paper are independent parameters. One cannot be predicted from knowing the other. Second, it shows that the difference in ar between the *fast-hi* and the *slow-late* clusters are independent of r. The quantity ar is a measure of how the rate increases upon citation for times scales less than $1/r$. Since $1/r$ is a characteristic time scale for the citation, we see in Figure 3a that the initial expected rate of citation is not a strong indicator of the longevity of a paper. Figure 3b shows the same separation from Figure 1 in terms of physical quantities, as ar is the initial increase in citation rate every time the paper is cited.

V. CAN WE PREDICT A FULL CITATION HISTORY OF A PAPER FROM ITS FIRST FEW YEARS?

How well can we estimate the three parameters of a paper's citation history, after observing only the first few years following publication? First, we consider a three-category classification problem. At a given year after publication, which we call the observation time, what is the fraction of papers more than 30 years old that are correctly classified into one of the three clusters using only the data from this training period? We note that for a 3-category classification, purely random chance will give us a 33% success rate. Figure 4 shows the prediction success as a function of training time. Success rises from 33%, reaches 50% at 5 years of training time, and continues increasing monotonically with increasing observation time.

Second, we also consider a 2-category classification problem. Here we seek to predict from the early citation history of a paper whether it will prove to be great (*fast-*

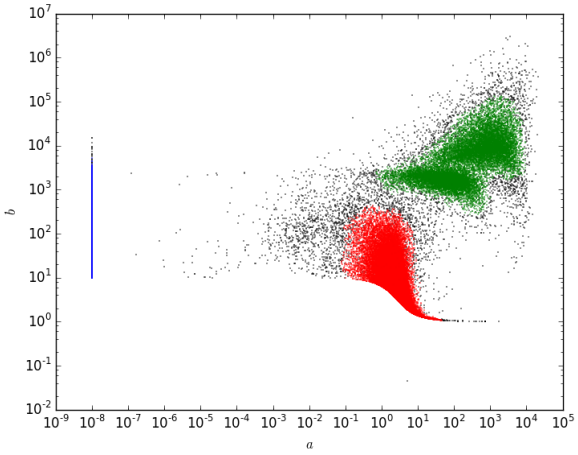

(a) b vs. a

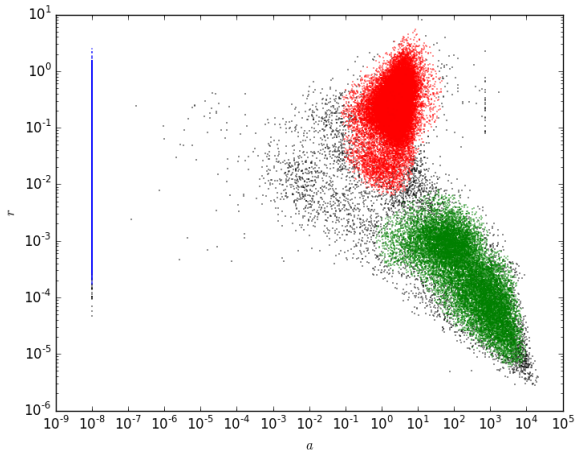

(b) r vs. a

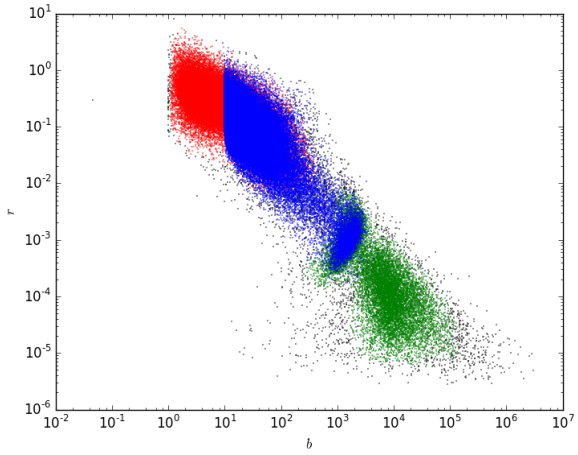

(c) r vs. b

FIG. 1: Scatter plot showing the clustering of papers based on the DBSCAN algorithm, with $\epsilon = 0.5$ and a minimum threshold of 50 papers projected onto (a) the b and a axes, (b) the r and a axes, and (c) the r and b axes. The different clusters are color coded as red (*fast-hi*), green (*slow-late*), and blue (*fast-flat*). The black points are papers categorized as noise. The parameter a is bounded from below by 10^{-8}.

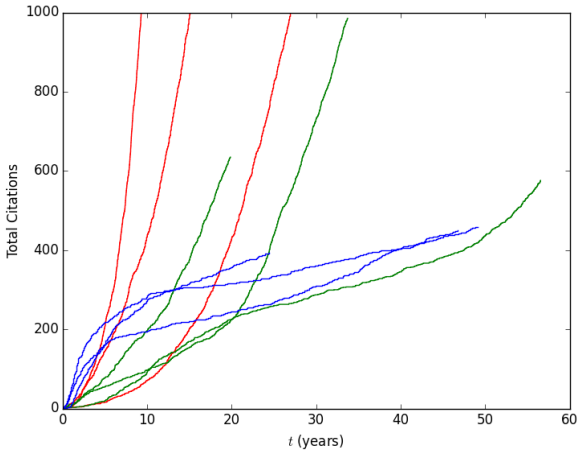

FIG. 2: Citation trajectories of the ten most cited papers in each cluster. The colors indicate the clusters the papers belong to and are the same as Figure 1

hi) or not. Random guessing would give a 50% success rate at this task. Figure 5 shows the 2-state classification success, precision and recall as a function of training time. Interestingly, there is a dip in the prediction success over the first few years, just because of the higher frequency of papers in the NOT FAST cluster and the very short training periods predicting the overwhelming majority are in the NOT FAST cluster (93.3% in the most extreme case). The figure shows that around half of the *fast-hi* papers can be identified with 16 to 17 years of training data (recall of 50%), a comparable precision of 50% would require 19 years of training data.

VI. CITERS ARE LOSING INTEREST IN PAPERS FASTER THAN BEFORE

The number of papers in the scientific literature increases exponentially[10]. A consequence is that interest in any given paper should naturally drop off faster in newer literature than in older literature[11]. We observe this phenomena in our analysis. We collected papers into age groups of 0 to 10, 10 to 20 and 30 to 100 years old. We compared the distribution of r for each of them; see Figure 6. We find that the older papers (lower panels) have minimal r values shifted to the left, indicating that citations diminish faster for newer papers than for older ones.

VII. THIS MODEL GIVES METRICS FOR COMPARING JOURNALS

We looked at the various numbers of papers in each category in different physics subjournals. We found an enrichment of *fast-hi* papers in Phys. Rev. B (PRB), Phys. Rev. Lett., and Rev. Mod. Phys. relative to the other journals. PRB has significantly higher enrichment in the number of FAST papers compared to PRA, PRC, PRD and PRE, which are not enriched (see appendix B). Figure 7 shows the fraction of papers in *fast-hi* cluster. Journals are order from highest to lowest impact factor (left to right) respectively. The fraction of *fast-hi* papers is significantly higher for Rev. Mod. Phys. compared to the other journals, and Phys. Rev. B and Phys. Rev. Letters are comparable. In this regard, our categorizations are consistent with PRL being high impact, RMP being a major review journal and PRB, a journal for condensed matter physics, which is the fastest growing subfield of physics in numbers of Ph.D.s awarded re-

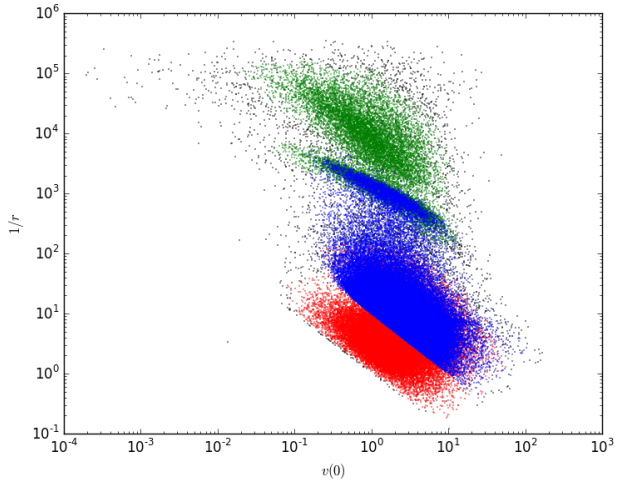

(a)

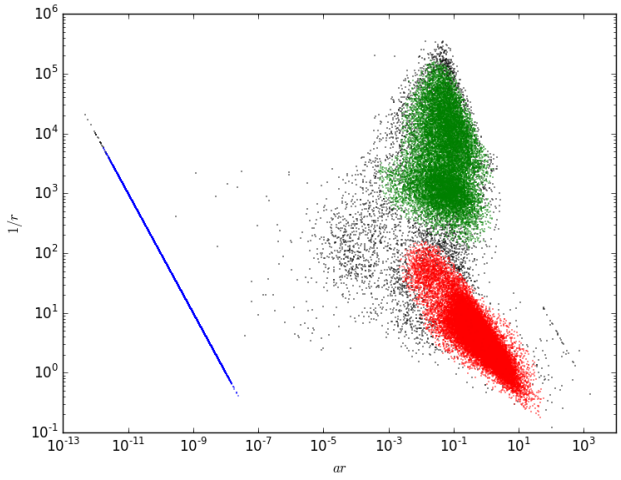

(b)

FIG. 3: We compare (a) the characteristic time scale $1/r$ versus expected initial velocity, and (b) the characteristic time scale $1/r$ versus the short term increase in citation rate per citation ar. The colors indicate the same clusters as Figure 1.

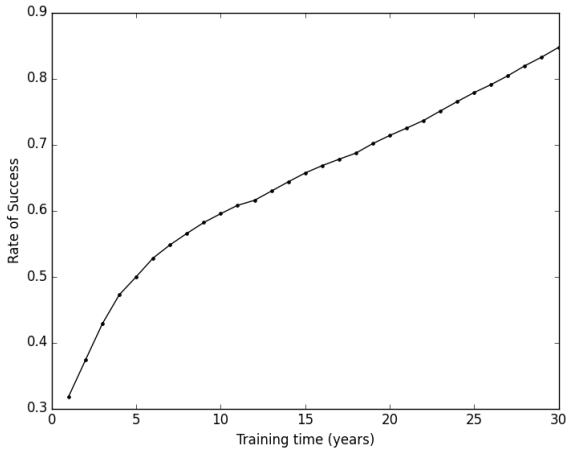

FIG. 4: Success Rate versus training time in years for the all classifiable papers over the age of 30 years.

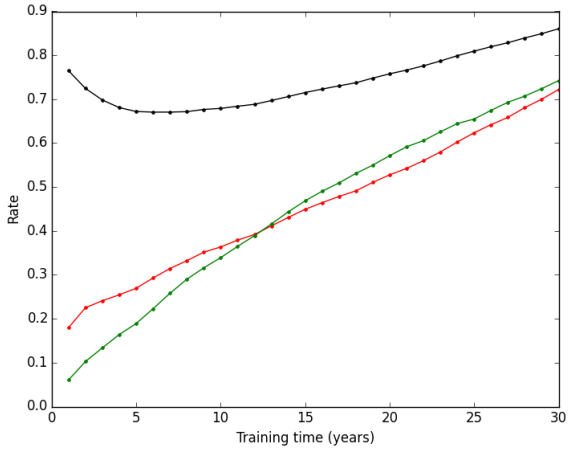

FIG. 5: Success Rate versus training time in years for the binary classification of *fast-hi* or not *fast-hi* in black, the precision (positive predictive value) in red and the recall (true positive rate) in green for all papers over 30 years of age.

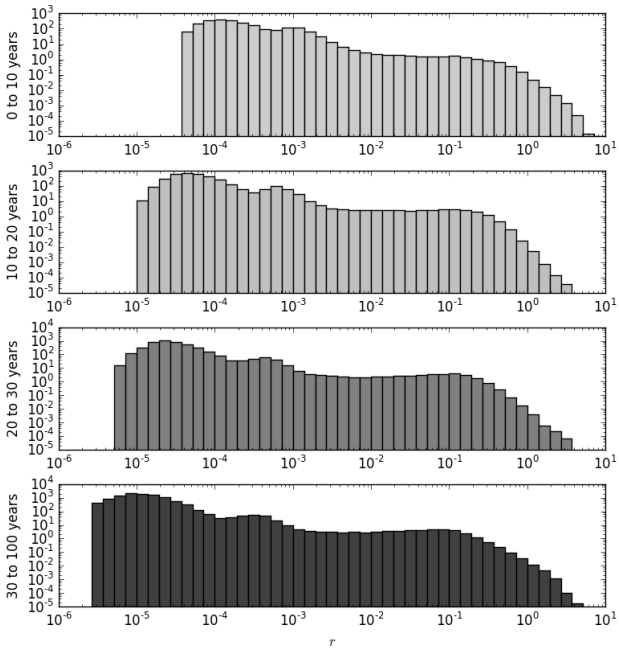

FIG. 6: The distribution of the decay parameter r for different age groups of papers. It is notable that the younger the papers are the more right-shifted the distribution of r.

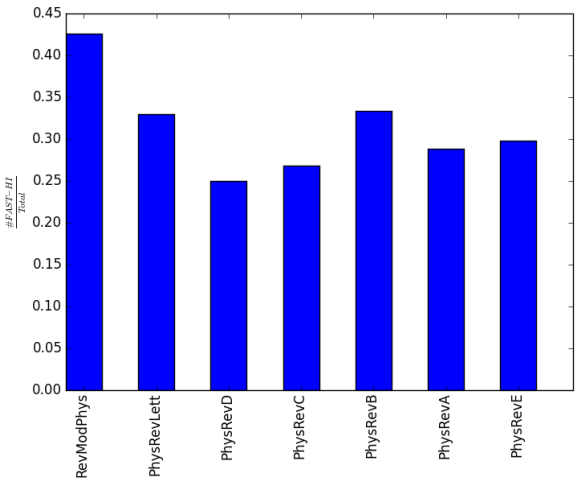

FIG. 7: A bar plot comparing the fraction of papers in the *fast-hi* cluster of the most important journals in Physical Review ordered by highest to lowest impact factor.

cently, nearly double that of its closest competitor particle physics[12–14].

VIII. CONCLUSION

Using a very simple model based on the pure-birth stochastic process, we characterize the typical behavior of papers, such as a "rich-get-richer" effect and a longer term decay in citation, without too many parameters. There is the behavior of sleepers which is not captured properly by the model, such as sleepers' have unreasonably large values of parameter a, but this does not prevent us from utilizing the model to categorize papers based on where they are located in parameter space. There are three categories which papers cluster into high impact (*fast-hi*), rapidly go obsolete (*fast-flat*) and late bloomers/sleeping beauties[15] (*slow-late*). We saw that Rev. Mod. Phys., Phys. Rev. Lett., and Phys. Rev. B have a higher proportion of *fast-hi* papers which correlates with the impact factor and yearly number of Ph.D.s

awarded. The categorization of papers becomes more accurate with the more years of data available, but after 5 years of training data more than half of the papers will be correctly identified.

[1] A. J. Lotka, Journal of the Washington Academy of Science **16**, 317 (1926).

[2] K. H. Chung and R. A. K. Cox, The Journal of Finance **45**, 301 (1990).

[3] M. E. J. Newman, "Power laws, Pareto distributions and Zipf's law," (2006), arXiv:0412044 [cond-mat].

[4] G. J. Peterson, S. Press, and K. A. Dill, Proceedings of the National Academy of Sciences **107**, 16023 (2010), http://www.pnas.org/content/107/37/16023.full.pdf.

[5] D. Wang, C. Song, and A.-L. Barabási, Science **342**, 127 (2013), http://science.sciencemag.org/content/342/6154/127.full.pdf.

[6] H.-W. Shen, D. Wang, C. Song, and A.-L. Barabási, "Modeling and predicting popularity dynamics via reinforced poisson processes," (2014), arXiv:1401.0778 [cs.SI].

[7] S. Xiao, J. Yan, C. Li, B. Jin, X. Wang, X. Yang, S. M. Chu, and H. Zha, in 25th International Joint Conference on Artificial Intelligence, edited by S. Kambhampati (AAAI Press / International Joint Conferences on Artificial Intelligence, Palo Alto, California USA, 2016) pp. 2676–2682.

[8] H. Panjer and G. Willmot, Insurance risk models (Society of Acturaries, 1992).

[9] P. O. Larsen and M. von Ins, Scientometrics **84**, 575 (2010).

[10] D. de Solla Price, Little science, big science, George B. Pegram lectures (Columbia University Press, 1963).

[11] P. D. B. Parolo, R. K. Pan, R. Ghosh, B. A. Huberman, K. Kaski, and S. Fortunato, J. Informetrics **9**, 734 (2015).

[12] P. J. Mulvey and S. Nicholson, "Physics Graduate Degrees: Results from the Enrollments and Degrees & the Degree Recipient Follow-up Surveys," https://tinyurl.com/zdz8kjn (2011).

[13] P. J. Mulvey and S. Nicholson, "Trends in Physics PhDs," https://tinyurl.com/hl4ytd2 (2014).

[14] "Physics PhDs granted by Subfield: Classes of 2013 & 2014 combined," https://tinyurl.com/j25qdh5 (2016).

[15] Q. Ke, E. Ferrara, F. Radicchi, and A. Flammini, Proceedings of the National Academy of Sciences **112**, 7426 (2015), http://www.pnas.org/content/112/24/7426.full.pdf.

Appendix A: Methods

We only included papers from the Physical Review corpus with more than 10 citations, to avoid overfitting and excluded problematic papers which were cited before publication. Given data containing the times t_1, t_2, t_3, ... , t_n for n events that occurs in a time interval from 0 to T such that $0 < t_1 < t_2 < t_3 < \cdots < t_n < T$ the model parameters can be extracted by maximizing the likelihood of observing the data given the model, without needing to repeat multiple trials. This is convenient, since we will be using the birth process to describe the citations of a paper and we cannot repeat history to get multiple trials, to maximize the probability $p_{m,n}(t,\tau)$. Knowing the probability of no events in a time interval τ given m events occurred in a time t,

$$p_{m,0}(t,\tau) = \exp\left(\int_0^\tau \lambda_m(t+u)\,du\right) \tag{A1}$$

allows us to determine the probability of any number of events occurring in the time interval τ given m events occurred in a time t, $p_{m,n>0}(t,\tau) = 1 - p_{m,0}(t,\tau)$. The probability $p_{m,n>0}(t,\tau)$ is the cumulative distribution function for the waiting time distribution, since it is the chance of observing at least one event over a time interval τ. Differentiating $p_{m,n>0}(t,\tau)$ with respect to τ will give the waiting time distribution or the probability that an event occurs after a time τ passes between τ and $\tau + d\tau$

$$f(\tau|m,t)d\tau = \lambda_m(t+\tau)\exp\left(-\int_0^\tau \lambda_m(t+u)\,du\right)d\tau \tag{A2}$$

so the probability density of n events occurring at the times t_1, t_2, t_3, ... , t_n with the condition that $0 < t_1 < t_2 < t_3 < \cdots < t_n < T$ is simply the products of the above formula which give the likelihood function

$$\begin{aligned}\mathcal{L}(a,b,\ldots|t_1,t_2,t_3,\ldots,t_n) &= f(t_1,t_2,t_3,\ldots,t_n|a,b,\ldots) \\ &= \exp\left(-\int_0^T \lambda_{N(t)}(t)\,dt\right)\prod_{i=1}^{n}\lambda_{i-1}(t_i)\end{aligned} \tag{A3}$$

Note that in the integral the rate $\lambda_{N(t)}(t)$ has an implicit time dependence on the number of events that happened up to the time t. Explicitly $N(t)$ can be written as

$$N(t) = \sum_{i=1}^{n}\theta(t-t_i) \tag{A4}$$

where $\theta(x)$ is the Heaviside step function which is unity for $x > 0$ and zero $x < 0$. The formula in equation (A3) is general, as we made no assumptions about the mathematical form of the rate. It is quite easy to interpret that maximum likelihood estimate (MLE) maximizes the rates at the points in time when events occur, but minimize the rate between the times of events.

The MLE fitting is done in Python using the minimize function in SciPy's optimize module. The optimization was given bounds and limited to positive values except for the parameter a which is bounded from below by 1×10^{-8} to avoid divide by zero errors. Initial guesses for the parameters were determined from the curve_fit function in SciPy's optimize module fitting the cumulative number of citations as a function of time with the expected number of citations given by equation (2) (for $m = 0$ and $t = 0$).

Appendix B: Enrichment of the clusters in the different journals

Of the papers fit to our statistical model and clustered we analyzed the composition of these journals in terms of the different clusters. We took the ratio of the fraction of papers in each cluster to be the fraction of the

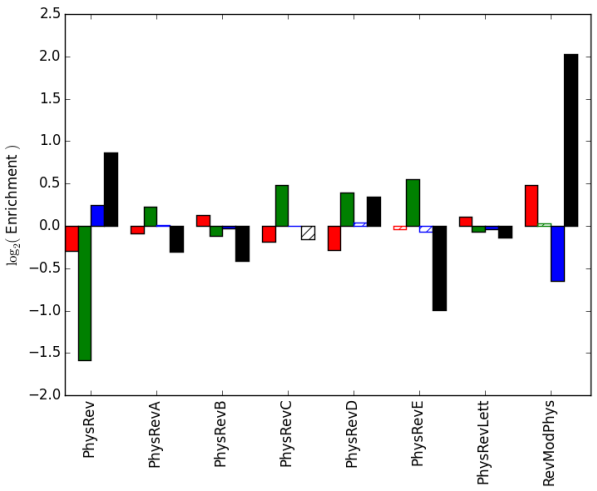

FIG. 8: Logarithm of Enrichment for each journal the colors correspond to those of figure 1. The solid bars are statistically significant at a level of $0.01/(4 \times 12) \approx 2.1 \times 10^{-4}$ the hatched bars are not significant at that level.

cluster on the whole. We call the ratio of these two fractions the enrichment. Mathematically this is expressed as $E_{c,j} = \frac{f_{c,j}}{f_{c,whole}}$ $f_{c,j}$ is the fraction of papers in cluster c belonging to the journal j and $f_{c,whole}$ is the fraction of papers in the cluster c on the whole of the Physical Review corpus. When the logarithm of the enrichment is positive (negative), it indicates a greater (lower) proportion of papers in that cluster for the respective journal compared to the whole. A value of zero means the two fractions are the same. A bar plot of the enrichments in Figure 8 shows that Phys. Rev. Lett., Phys. Rev. B and Rev. Mod. Phys are the only journals to have a positive enrichment in the *fast-hi* papers negative enrichment in papers from the *fast-flat* cluster.

Journal	whole	*fast-hi*	*fast-hi* p-value	LATE	LATE p-value	*fast-flat*	*fast-flat* p-value	Noise	Noise p-value
PhysRev	13682	3406	2.4e-31	475	3.2e-187	9142	1.1e-40	659	1.4e-48
PhysRevA	15309	4406	7.1e-06	1866	2.9e-12	8708	6.7e-03	329	5.4e-06
PhysRevB	39985	13326	1.7e-24	3844	1.6e-09	22020	4.1e-05	795	8.6e-23
PhysRevC	7877	2114	2.5e-09	1149	1.0e-26	4427	1.2e-02	187	9.2e-03
PhysRevD	18361	4595	1.5e-40	2509	7.3e-41	10640	3.9e-04	617	2.8e-10
PhysRevE	6010	1792	8.8e-03	915	3.1e-26	3223	9.5e-04	80	1.1e-12
PhysRevLett	48246	15922	6.7e-25	4782	3.0e-06	26376	7.1e-08	1166	1.0e-05
PhysRevSTAB	99	24	6.5e-02	30	2.5e-06	41	2.5e-02	4	1.5e-01
PhysRevSTPER	10	1	1.9e-01	5	1.2e-02	3	2.0e-01	1	2.4e-01
PhysRevSeriesI	5	1	4.3e-01	0	6.1e-01	2	3.9e-01	2	1.6e-02
PhysRevX	20	7	2.2e-01	3	2.3e-01	5	5.2e-02	5	4.5e-04
RevModPhys	1478	630	1.8e-12	157	3.6e-02	531	9.4e-21	160	3.4e-46

TABLE I: The number of papers in each cluster and journal compared to the whole with p-value determined by a one-sided Fisher exact test.